Stocks as Lotteries: The Implications of Probability Weighting for Security Prices

Nicholas Barberis and Ming Huang
Yale University and Stanford / Cheung Kong University

September 2004*

Abstract

As part of their cumulative prospect theory (CPT), Tversky and Kahneman (1992) argue that, when people evaluate risk, they transform objective probabilities via a weighting function which, among other features, overweights small probabilities. We investigate the implications of CPT for the pricing of financial securities, paying particular attention to the effects of the weighting function. Under CPT, the CAPM can hold when securities are normally distributed; but a positively skewed security can become overpriced and earn very low average returns, even if small and independent of other risks, and even if just one of many skewed securities in the economy. We apply the last result to the pricing of IPOs and to the valuation of equity stubs. Using data on the skewness of IPO returns, we show that investors with CPT preferences calibrated to experimental evidence would require an average return on IPOs that is several percentage points below the market return. Under CPT, then, the historical underperformance of IPOs may not be so puzzling.

*Preliminary and incomplete. Please do not distribute without permission. We are grateful to Tongshu Ma, Jay Ritter, and Joshua Schaeffer for assistance with data. Comments welcome at nick.barberis@yale.edu and mhuang@ckgsb.edu.cn.

1 Introduction

Over the past few years, researchers have accumulated a large body of experimental evidence on attitudes to risk. As is well known, one of the findings of this literature is that, when people evaluate risk, they routinely violate the predictions of expected utility. Of course, the fact that people depart from expected utility in experimental settings does not necessarily mean that they also do so in financial markets. Nonetheless, given the difficulties the expected utility framework has encountered in addressing a number of financial phenomena, it may be useful to check whether incorporating some of the experimental findings into our models of investor behavior can improve these models' ability to match the empirical facts.

Tversky and Kahneman's (1992) "cumulative prospect theory," a modified version of "prospect theory" (Kahneman and Tversky, 1979), is perhaps the best available summary of typical attitudes to risk in experimental settings. Under this theory, people evaluate risk using a value function that is defined over gains and losses, is concave over gains and convex over losses, and is kinked at the origin; and using *transformed* rather than objective probabilities, where the transformed probabilities are obtained from objective probabilities by applying a weighting function. A key aspect of this weighting function is the overweighting of small probabilities – a feature inferred from people's demand for lotteries offering a small chance of a large gain, and for insurance protecting against a small chance of a large loss.

In this paper, we study the pricing of financial securities when investors make decisions according to cumulative prospect theory. Previous research on this topic has focused mainly on the implications of the kink in the value function (Benartzi and Thaler, 1995, Barberis, Huang and Santos, 2001). Here, we turn our attention to other, less-studied aspects of cumulative prospect theory and, in particular, to the probability weighting function.

First, we show that in a one-period equilibrium setting with normally distributed security payoffs and homogeneous investors, the CAPM can hold even when investors evaluate risk according to cumulative prospect theory. The intuition is that, since cumulative prospect theory preferences do satisfy first-order stochastic dominance, investors still want to hold portfolios on the mean-variance efficient frontier; in other words, portfolios that are some combination of the risk-free asset and the tangency portfolio. Market clearing means that the tangency portfolio must be the market portfolio, and the CAPM follows in the usual way.

We then show, however, that if a non-normally distributed security – specifically, a positively skewed security – is introduced into the economy, probability weighting can have unusual pricing effects. In particular, the security can become overpriced, relative to the price that would be set by investors who do not weight probabilities, and can earn a very low average return, even if it is small relative to total market capitalization and independent of other risks. Intuitively, if the new security is sufficiently skewed, some investors may choose to hold undiversified portfolios that take large positions in it, thereby making the distribution of their overall wealth more lottery-like. Since a cumulative prospect theory investor overweights small probabilities, he loves lottery-like wealth distributions, and is therefore willing to pay a very high price for the skewed security. In deriving this result, we assume the existence of short-sale constraints; just how overpriced the skewed security is, therefore depends on how high short-selling costs are.

Our result that a skewed security can become overpriced is by no means an obvious one. An investor who overweights low probabilities will, of course, value a skewed *portfolio* highly; but what is surprising is that he also values a skewed *security* highly, even if that security is small and independent of other risks. Indeed, as we discuss below, such a security would *not* earn a low average return in a traditional expected utility model of skewness preference.

Our analysis also reveals some other implications of probability weighting for pricing. We show, for example, that under cumulative prospect theory, the relationship between a security's skewness and its expected return is very nonlinear: a highly skewed security can be overpriced and earn a low average return, but a security that is merely moderately skewed is priced fairly. We also show that a positively skewed security can be overpriced even if there are *many* skewed securities in the economy. Intuitively, even with a number of skewed securities at their disposal, investors may prefer a large undiversified position in just one skewed security to a diversified position in several.

This last result offers a new way of thinking about a number of puzzling asset market phenomena, most notably the low average return of IPOs (Ritter, 1991).[1] IPOs have a positively skewed return distribution, probably because, being young firms, a large fraction of their value is in the form of growth options. Our analysis implies that, in an economy populated by investors with cumulative prospect theory preferences, IPOs can become overpriced and earn low average returns, so long as their skewness is sufficiently high. We use Ritter's (1991) original data to measure the level of skewness in IPO returns. According to our model, for this level of skewness, investors with cumulative prospect theory preferences calibrated to experimental data would require an average return that is several percentage points below the market return. Under cumulative prospect theory, then, the historical performance of IPOs may not be so puzzling.

Our results may also shed light on so-called "equity stub" anomalies, whereby parent companies appear undervalued relative to their publicly traded subsidiaries (Mitchell, Pulvino, and Stafford, 2002, Lamont and Thaler, 2003). If a firm has a publicly traded subsidiary that is valued mainly for its growth options, that subsidiary's stock returns are likely to be

[1] Just how low these average returns are, is a matter of debate. Most recently, researchers have investigated the extent to which the low average returns may reflect "pseudo market timing" (Schultz, 2003, Dahlquist and De Jong, 2004).

positively skewed. Our analysis suggests that investors will then overprice the subsidiary relative to the parent company, a prediction that is at least qualitatively consistent with the empirical facts.

The potential application of cumulative prospect theory to the overpricing of IPOs has also been noted by Brav and Heaton (1996) and Ma and Shen (2003). However, in writing down the equilibrium conditions for their quantitative analysis, these papers assume that investors engage in "narrow framing" – in other words, that investors get utility *directly* from IPO returns, even if IPO stocks are only a small part of their portfolios. While narrow framing may well occur in practice, we believe that, as a first step, it is better to analyze the implications of probability weighting under the traditional assumption that investors get utility only from overall *portfolio* returns, not from individual security returns.

Through the probability weighting function, cumulative prospect theory investors exhibit a preference for skewness. There are already a number of papers that analyze the implications of skewness-loving preferences (Kraus and Litzenberger, 1976). We note, however, that the pricing effects we demonstrate are new to the skewness literature. Earlier papers have shown that, in the presence of skewed securities, coskewness with the market can earn a risk premium. We show that skewness itself, and not just coskewness with the market, can earn a premium. For example, in our economy, a skewed security can earn a negative average return even if it is small and independent of other risks; in other words, even if its coskewness with the market is zero.

How do we obtain this new effect? The earlier skewness literature only considers economies in which all investors hold diversified portfolios. In such economies, only coskewness with the market portfolio earns a premium. We show that, under cumulative prospect theory, there can be equilibria in which the presence of a skewed security induces some agents to hold undiversified portfolios. In such economies, it is not only coskewness with the market that matters, but also a security's skewness itself.

Our more surprising result – that, under cumulative prospect theory, a skewed security can be overpriced even if small and independent of other risks, and even if just one of many skewed securities in the economy – is, to our knowledge, completely new to the literature. The fact that the CAPM can hold even under cumulative prospect theory has been recently argued by Levy, De Giorgi and Hens (2003). On this point, our contribution is to offer a very different proof of the result.

In Section 2, we discuss cumulative prospect theory and its probability weighting feature in more detail. In Section 3, we present our assumptions on investors' preferences. We then examine how these investors price normally distributed securities (Section 4) and positively skewed securities (Section 5). Section 6 considers applications of our results, with particular emphasis on the pricing of IPOs. Section 7 concludes.

2 Cumulative Prospect Theory and Probability Weighting

We introduce Tversky and Kahneman's (1992) cumulative prospect theory by first reviewing the original prospect theory, laid out by Kahneman and Tversky (1979), on which it is based.

Consider a gamble

$$(x, p; y, q),$$

to be read as "get x with probability p and y with probability q, independent of other risks," where $x < 0 < y$ or $y < 0 < x$, and where $p + q = 1$. In the expected utility framework, the agent evaluates this risk by computing

$$pu(W + x) + qu(W + y), \tag{1}$$

where W is his current wealth. In the original version of prospect theory, presented in Kahneman and Tversky (1979), the agent assigns the gamble the value

$$\pi(p)v(x) + \pi(q)v(y), \tag{2}$$

where $v(\cdot)$ and $\pi(\cdot)$ are shown in Figure 1.

Four differences between (1) and (2) should be noted. First, the carriers of value in prospect theory are gains and losses, not final wealth positions: the argument of $v(\cdot)$ in (2) is x, not $W + x$. This is motivated in part by explicit experimental evidence, and in part by introspection: we often perceive the level of attributes – brightness, loudness or temperature, say – relative to their earlier levels, rather than in absolute terms.

Second, the value function $v(\cdot)$ is concave over gains, but convex over losses. This is inferred from subjects' preference for a certain gain of \$500 over[2]

$$(\$1,000, \frac{1}{2}),$$

and from their preference for

$$(-\$1,000, \frac{1}{2})$$

over a certain loss of \$500. In short, people are risk averse over moderate-probability gains, but risk-seeking over moderate-probability losses.

Third, the value function is kinked at the origin, so that the agent is more sensitive to losses – even small losses – than to gains of the same magnitude. Kahneman and Tversky (1979) infer the kink from the widespread aversion to bets of the form

$$(\$110, \frac{1}{2}; -\$100, \frac{1}{2}).$$

[2]We abbreviate $(x, p; 0, q)$ to (x, p).

Such aversion is hard to explain with differentiable utility functions, whether expected utility or non-expected utility, because the very high local risk aversion required to do so typically predicts implausibly high aversion to large-scale gambles (Epstein and Zin, 1990, Rabin, 2000, Barberis, Huang and Thaler, 2003).

Finally, under prospect theory, the agent does not use objective probabilities when evaluating the gamble, but rather, transformed probabilities obtained from objective probabilities via a probability weighting function $\pi(\cdot)$. This function has two salient features. First, small probabilities are overweighted: in the lower panel of Figure 1, the solid line lies above the dotted line for low p. Given the concavity (convexity) of the value function in the region of gains (losses), this is inferred from people's preference for

$$(\$5,000, 0.001)$$

over a certain \$5, and from their preference for a certain loss of \$5 over

$$(-\$5,000, 0.001);$$

in other words, from their simultaneous demand for both lotteries and insurance.

The other main feature of the probability weighting function is a greater sensitivity to differences in probability at higher probability levels: in the lower panel of Figure 1, the solid line is flatter for low p than for high p. For example, subjects tend to prefer a certain \$3,000 to (\$4,000, 0.8), but also prefer (\$4,000, 0.2) to (\$3,000, 0.25). This pair of choices violates expected utility, but under prospect theory, implies

$$\frac{\pi(0.25)}{\pi(0.2)} < \frac{\pi(1)}{\pi(0.8)}. \tag{3}$$

The intuition is that the 20 percent jump in probability from 0.8 to 1 is more striking to people than the 20 percent jump from 0.2 to 0.25. In particular, people place much more weight on outcomes that are certain relative to outcomes that are merely probable, a feature sometimes known as the certainty effect.

Tversky and Kahneman (1992) present a modified version of prospect theory, known as cumulative prospect theory. In this version, they provide explicit functional forms for $v(\cdot)$ and $\pi(\cdot)$. Moreover, they apply the probability weighting function to the cumulative probability distribution, not to the probability density function. This ensures that the preferences do not violate first-order stochastic dominance – a weakness of the original 1979 version of prospect theory – and also that they can be applied to gambles with any number of outcomes, not just two. Finally, Tversky and Kahneman (1992) allow the probability weighting functions for gains and losses to be different.

In full, cumulative prospect theory says that the agent evaluates a gamble

$$(x_{-m}, p_{-m}; \ldots; x_{-1}, p_{-1}; x_0, p_0; x_1, p_1; \ldots; x_n, p_n),$$

where $x_i < x_j$ for $i < j$ and $x_0 = 0$, by assigning it the value

$$\sum_{i=-m}^{n} \pi_i v(x_i), \tag{4}$$

where

$$\pi_i = \left\{ \begin{array}{ll} w^+(p_i + \ldots + p_n) - w^+(p_{i+1} + \ldots + p_n) & \text{for} \quad 0 \leq i \leq n \\ w^-(p_{-m} + \ldots + p_i) - w^-(p_{-m} + \ldots + p_{i-1}) & \phantom{\text{for}} \quad -m \leq i < 0 \end{array} \right. , \tag{5}$$

and where $w^+(\cdot)$ and $w^-(\cdot)$ are the probability weighting functions for gains and losses, respectively.

Equation (5) emphasizes that under cumulative prospect theory, the weighting function is applied to the cumulative probability distribution. If it were instead applied to the probability density function, as it is in the original prospect theory, the probability weight π_i, for $i \geq 0$ say, would be $w^+(p_i)$, which can be trivially rewritten as

$$w^+([p_i + p_{i+1} + \ldots + p_n] - [p_{i+1} + \ldots + p_n]). \tag{6}$$

Comparing this with equation (5), we see that under cumulative prospect theory, the probability weight π_i is obtained by first applying the weighting function to the cumulative probability distribution and then taking differences, while in (6), the order is reversed: we first take differences and then apply the weighting function.

Tversky and Kahneman (1992) propose the functional forms

$$v(x) = \left\{ \begin{array}{ll} x^\alpha & \text{for} \quad x \geq 0 \\ -\lambda(-x)^\alpha & \phantom{\text{for}} \quad x < 0 \end{array} \right. , \tag{7}$$

and

$$w^+(P) = w^-(P) = w(P) = \frac{P^\delta}{(P^\delta + (1-P)^\delta)^{1/\delta}}. \tag{8}$$

For $0 < \alpha < 1$ and $\lambda > 1$, $v(\cdot)$ captures the features of the value function highlighted earlier: it is concave over gains, convex over losses, and kinked at the origin. The severity of the kink is determined by λ, which measures the relative sensitivity to gains and losses and is known as the coefficient of loss aversion. For $0 < \delta < 1$, $w(\cdot)$ captures the features of the weighting function described earlier: it overweights low probabilities, so that $w(P) > P$ for low P, and it is flatter for low P than for high P.

Based on their experimental evidence, Tversky and Kahneman (1992) estimate $\alpha = 0.88$, $\lambda = 2.25$ and $\delta = 0.65$. Figure 2 shows the form of the probability weighting function $w(\cdot)$ for $\delta = 0.65$ (the dashed line), for $\delta = 0.4$ (the dash-dot line), and for $\delta = 1$, which corresponds to no probability weighting at all (the solid line). The overweighting of small probabilities and the greater sensitivity to changes in probability at higher probability levels are both clearly visible for $\delta < 1$.

In our subsequent analysis, we work with cumulative prospect theory – in other words, with the specification in equations (4)-(5) and (7)-(8), adjusted only to allow for continuous probability distributions.

3 Investor Preferences

In Sections 4 and 5, we analyze security pricing in an economy where investors use cumulative prospect theory to evaluate gambles – and who therefore, in particular, transform probabilities according to a weighting function. In this section, we start by specifying investor preferences.

Each agent has beginning-of-period wealth W_0 and end-of-period wealth $\widetilde{W} = W_0 \widetilde{R}$. Since prospect theory defines utility over gains and losses measured relative to a reference level, we introduce a reference wealth level W_z, and define utility over

$$\widehat{W} = \widetilde{W} - W_z, \tag{9}$$

or equivalently, over

$$\widehat{R} = \widetilde{R} - R_z, \tag{10}$$

where $R_z = W_z / W_0$. In words, a wealth outcome that exceeds (falls below) W_z is labelled a gain (loss). One possible value of W_z is W_0. Another, that we adopt in this paper, is $W_0 R_f$, where R_f is the gross risk-free rate, so that

$$\widehat{W} = \widetilde{W} - W_0 R_f, \quad \widehat{R} = \widetilde{R} - R_f.$$

In this case, the agent only considers the return on his wealth a gain if it exceeds the risk-free rate.

We further assume:

Assumption 1: $|E(\hat{W})|$, $\mathrm{Var}(\hat{W}) < \infty$, or equivalently, $|E(\hat{R})|$, $\mathrm{Var}(\hat{R}) < \infty$.

The investor's goal function is:

$$U(\widetilde{W}) \equiv V(\hat{W}) = V(\hat{W}^+) + V(\hat{W}^-), \tag{11}$$

where

$$\begin{aligned} V(\hat{W}^+) &= -\int_0^\infty v(W)\, dw^+(1 - P(W)) && (12) \\ V(\hat{W}^-) &= \int_{-\infty}^0 v(W)\, dw^-(P(W)), && (13) \end{aligned}$$

and where $P(\cdot)$ is the cumulative probability distribution function. Equations (11)-(13) correspond directly to equations (4)-(5), modified to allow for continuous probability distributions. As before, $w^+(\cdot)$ and $w^-(\cdot)$ are the probability weighting functions for gains and losses, respectively. We assume:

Assumption 2: $w^+(\cdot) = w^-(\cdot) \equiv w(\cdot)$.

Assumption 3: $w(\cdot)$ takes the form proposed by Tversky and Kahneman (1992),

$$w(P) = \frac{P^\delta}{(P^\delta + (1-P)^\delta)^{1/\delta}}, \tag{14}$$

where $\delta \in (0,1)$. As mentioned above, experimental evidence suggests $\delta \approx 0.65$.

Assumption 4: $v(\cdot)$ takes the form proposed by Tversky and Kahneman (1992),

$$v(x) = \begin{cases} x^\alpha & \text{for } x \geq 0 \\ -\lambda(-x)^\alpha & \text{for } x < 0, \end{cases} \tag{15}$$

where $\lambda > 1$ and $\alpha \in (0,1)$. As noted earlier, experimental evidence suggests $\alpha \approx 0.88$ and $\lambda \approx 2.25$.

Assumption 5: $\alpha < 2\delta$.

When taken together with Assumptions 1-4, Assumption 5 ensures that the goal function $V(\cdot)$ in (11) is well-behaved at $\pm\infty$, and therefore well-defined. The values of α and δ estimated by Tversky and Kahneman (1992) satisfy this condition. In the case of normal or lognormal distributions, Assumption 5 is not needed.

Before embarking on our formal analysis, we present a useful lemma.

Lemma 1: *Under Assumptions 1-5,*

$$V(\hat{W}^+) = \int_0^\infty w(1-P(x))dv(x) \tag{16}$$

$$V(\hat{W}^-) = -\int_{-\infty}^0 w(P(x))dv(x). \tag{17}$$

Proof of Lemma 1: See the Appendix for a full derivation. In brief, the lemma follows by applying integration by parts to equations (12) and (13).

4 The Pricing of Normally Distributed Securities

We now turn to the implications of cumulative prospect theory for asset pricing. We first show that, in a simple one-period equilibrium with normally distributed security payoffs, the CAPM can hold even when investors make decisions according to cumulative prospect theory.

Formally, suppose that there are J non-degenerate risky securities, where security j earns a gross return $\widetilde{R}_j$, and where

$$\widetilde{R} = (\widetilde{R}_1, \widetilde{R}_2, ..., \widetilde{R}_J).$$

If there is a risk-free asset, its gross return is denoted R_f.

We make the following assumptions:

Assumption 6: One time period. Agents derive utility only from time 1 consumption.[3]

Assumption 7: Identical preferences. All agents have the same preferences over time 1 consumption, namely those described in equations (11)–(13) and Assumptions 2-5.

Assumption 8: Supply of securities. Security j has $n_j > 0$ shares outstanding and a per-share payoff of $\tilde{X}_j$ at time 1. The J random payoffs $\{\tilde{X}_1, \cdots, \tilde{X}_J\}$ have a positive-definite variance-covariance matrix, so that no linear combination of the J payoffs is a constant.

Assumption 9: Homogeneous expectations. All agents assign the same probability distribution to future payoffs and security returns.

Assumption 10: Multivariate normality. The joint distribution of time 1 security payoffs is multivariate normal.

Assumption 11: All endowments are traded.

Assumption 12: There are no trading frictions or constraints.

We can now prove:

Proposition 1: *Suppose that Assumptions 6-12 hold, that there is a risk-free asset, and that the reference return $R_z = R_f$. Then there exists an equilibrium in which the CAPM holds, so that*

$$E(\tilde{R}_i) - R_f = \beta_i(E(\tilde{R}_M) - R_f), \quad i = 1, \cdots, n, \tag{18}$$

where

$$\beta_i \equiv \frac{\mathrm{Cov}(\tilde{R}_i, \tilde{R}_M)}{\mathrm{Var}(\tilde{R}_M)}. \tag{19}$$

[3] The analysis can easily be extended to the case of two periods, with agents also deriving utility from time 0 consumption.

The excess market return $\hat{R}_M \equiv \tilde{R}_M - R_f$ *satisfies*

$$V(\hat{R}_M) \equiv -\int_{-\infty}^{0} w(P(\hat{R}_M))dv(\hat{R}_M) + \int_{0}^{\infty} w(1 - P(\hat{R}_M))dv(\hat{R}_M) = 0 \tag{20}$$

and the market risk premium is positive.

Proof of Proposition 1: See the Appendix.

The intuition behind Proposition 1 becomes much clearer once we state the following proposition, which establishes that the preferences in (11)-(13) satisfy first-order stochastic dominance.

Proposition 2: *Under Assumptions 1-5, the preferences in (11)–(13) satisfy the first-order stochastic dominance (FOSD) property. That is, if* $\hat{W}_1$ *(strictly) first-order stochastically dominates* $\hat{W}_2$*, then* $V(\hat{W}_1) \geq (>) V(\hat{W}_2)$.[4]

Proof of Proposition 2: Since $\hat{W}_1$ FOSD $\hat{W}_2$, we have $P_1(x) \leq P_2(x)$ for all $x \in R$. Equations (16) and (17) immediately imply $V(\hat{W}_1^+) \geq V(\hat{W}_2^+)$ and $V(\hat{W}_1^-) \geq V(\hat{W}_2^-)$, and therefore that $V(\hat{W}_1) \geq V(\hat{W}_2)$. If, moreover, $\hat{W}_1$ strictly FOSD $\hat{W}_2$, then $P_1(x) < P_2(x)$ for some $x \in R$. Given that distribution functions are all right continuous, we have $V(\hat{W}_1) > V(\hat{W}_2)$.[5]

Proposition 1 now follows very easily. With normally distributed security payoffs, the goal function in (11)-(13) becomes a function of the mean and standard deviation of the distribution of overall wealth. Since the FOSD property holds for these preferences, all investors choose a portfolio on the mean-variance efficient frontier, in other words, a portfolio that is some combination of the risk-free asset and the tangency portfolio. Market clearing means that the tangency portfolio must be the market portfolio, and the CAPM follows in the usual way.

5 The Pricing of Positively Skewed Securities

Suppose that, in addition to the risk-free asset and the J normally distributed risky assets of Section 4, we also introduce a positively skewed security with gross return $\widetilde{R}_n \equiv \widehat{R}_n + R_f$.

[4]Proposition 2 is hardly surprising: after all, Tversky and Kahneman (1992) themselves point to the FOSD property as one of the advantages of cumulative prospect theory over prospect theory (see also Levy, De Giorgi, and Hens, 2003).

[5]Not all of Assumptions 1-5 are needed for this result. Assumptions 2-4 can be replaced by "$w(\cdot)$ and $v(\cdot)$ are strictly increasing and continuous," and Assumption 5 can be replaced by "the integrals in equations (12) and (13) are both well-defined and finite."

We now show that the preferences in (11)-(13) can induce unusual effects on the price of this security; in particular, the security can become overpriced and earn a low average return.

For the remainder of the analysis, we make the following simplifying assumptions:

Assumption 13: Independence. The new security's return is independent of the returns of the J existing risky securities.

Assumption 14: Small Supply. The total payoff or market value of the new security is small compared to the total payoff or market value of the J existing securities.

In a representative agent economy with concave, expected utility preferences, these two assumptions would imply a zero risk premium for the new security.

We now show that, under the preferences in (11)-(13), we can construct an equilibrium in which the new security earns a negative risk premium. In this equilibrium, there is heterogeneity in the portfolios that agents choose to hold: some agents hold no position at all in the new security, while others hold large positions in it. We achieve this heterogeneity not through heterogeneity in preferences – on the contrary, we assume that all agents have the same preferences – but rather, from the existence of non-unique optimal portfolios.

More specifically, we suppose that there are short-sales constraints, and show that there is an equilibrium in which holding a zero position in the new security and holding a position $x^* > 0$ in it are both global optima. To derive the equilibrium conditions, note first that the optimal portfolios will necessarily be some combination of R_f, $\widetilde{R}_M$ and $\widetilde{R}_n$, where $\widetilde{R}_M$ is the return on the market portfolio *excluding* the new security. The reason is that, even if the agent holds a position in the new security, the optimal combination of the original $J+1$ securities must still have the highest mean for given standard deviation, and must therefore still be on the mean-variance efficient frontier.[6]

The equilibrium conditions are therefore:

$$\begin{aligned} V(\widehat{R}_M) &= V(\widehat{R}_M + x^*\widehat{R}_n) = 0 && (21) \\ V(\widehat{R}_M + x\widehat{R}_n) &< 0 \text{ for } 0 < x \neq x^*, && (22) \end{aligned}$$

where

$$V(\hat{R}_M + x\hat{R}_n) = -\int_{-\infty}^{0} w(P_x(R))dv(R) + \int_{0}^{\infty} w(1 - P_x(R))dv(R) \qquad (23)$$

[6]Short-sale constraints are not crucial to our results. In the absence of such constraints, we can again construct an equilibrium with two global optima, one involving a substantial short position and the other, a substantial long position. While such an equilibrium generates similar results to the one we study, it also predicts that the aggregate long and aggregate short positions in the new security are both much larger than the security's total supply, which does not seem realistic.

and

$$P_x(R) = \Pr(\widehat{R}_M + x\widehat{R}_n \leq R). \tag{24}$$

In words, if x is the investor's allocation to the new security, then utility is maximized either by investing nothing at all in the new security ($x = 0$), or by investing a substantial positive amount in it ($x = x^*$). Any other positive allocation x yields lower utility.

5.1 Example

We now construct an explicit equilibrium in which conditions (21)-(22) hold. Given Assumption 10, the market return - again, excluding the new security – is normally distributed:

$$\widehat{R}_M \sim N(\mu_M, \sigma_M). \tag{25}$$

We model the positive skewness of the new security in the simplest possible way, using a binomial distribution:

$$\widehat{R}_n \sim (L, q; -R_f, 1-q), \tag{26}$$

where L is large and positive and where q is small. With high probability, the security pays out nothing, so that its gross return $\widetilde{R}_n$ is zero, and $\widehat{R}_n = \widetilde{R}_n - R_f = -R_f$, where, as in Section 4, the reference return R_z is taken to be the risk-free rate. With low probability, the security delivers a large positive return $\widehat{R}_n = L$. For *very* large L and *very* low q, the security's return distribution resembles that of a lottery ticket.

We take the unit of time to be a year. To start, we set the standard deviation of the market return, σ_M, to 0.15, and solve equation (20) for the market risk premium μ_M. In our benchmark case, the preference parameters take the values estimated by Tversky and Kahneman (1992), namely

$$\alpha = 0.88,\ \delta = 0.65,\ \lambda = 2.25.$$

For these parameter values, we find $\mu_M = 7.47$ percent. This is consistent with the work of Benartzi and Thaler (1995) and Barberis, Huang and Santos (2001), who show that prospect theory can generate a very substantial equity premium. The intuition is that since, by virtue of the kink in $v(\cdot)$, the investor is much more sensitive to losses than to gains, he perceives the stock market to be very risky, and charges a high average return as compensation.

We now set $L = 10$, which, for low q, implies a substantial amount of positive skewness in the new security's return $\widehat{R}_n$. For this L, for the assumed $\sigma_M = 0.15$, for the implied $\mu_M = 0.0747$, and for a gross risk-free rate $R_f = 1.02$, we search for q such that conditions (21)-(22) hold. Given our assumptions about the distribution of security returns, $P_x(R)$, defined in (24), is given by

$$
\begin{aligned}
P_x(R) &= \Pr(\widehat{R}_M + x\widehat{R}_n \leq R) \\
&= \Pr(\widehat{R}_n = L)\Pr(\widehat{R}_M \leq R - xL) + \Pr(\widehat{R}_n = -R_f)\Pr(\widehat{R}_M \leq R + xR_f) \\
&= qN(\frac{R - xL - \mu_M}{\sigma_M}) + (1-q)N(\frac{R + xR_f - \mu_M}{\sigma_M}),
\end{aligned} \tag{27}
$$

where $N(\cdot)$ is the cumulative Normal distribution.

We find that $q = 0.087$ satisfies conditions (21)-(22). The solid line in Figure 3 plots the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ for this q over a range of values of x, the allocation to the new security. The graph clearly shows the two global optima at $x = 0$ and $x = 0.084$.

The two optima in Figure 3 are trading off skewness and diversification. By taking a substantial position of $x = 0.084$ in the skewed security, the investor, on the one hand, gets something that he values highly – a small chance of becoming very wealthy – but, on the other hand, leaves himself severely undiversified. If the new security's skewness is just strong enough to compensate for the lack of diversification, there are two global optima.

The final step is to compute the new security's equilibrium expected return. Given the distribution in (26), this is given by

$$E(\hat{R}_n) = qL - (1-q)R_f. \tag{28}$$

In our example, $E(\widehat{R}_n) = -0.0723$, which, given that $R_f = 1.02$, implies a gross expected return on the skewed security of $E(\widetilde{R}_n) = E(\widehat{R}_n) + R_f = 0.9477$, and hence a net expected return of -5.23 percent.

We have therefore shown that, under cumulative prospect theory, a positively skewed security can become overpriced and can earn low average returns, even if it is small relative to total market capitalization and independent of other risks. Intuitively, if the new security is sufficiently skewed, some investors may choose to hold undiversified portfolios that take large positions in it, thereby making the distribution of their overall wealth more lottery-like. Since a cumulative prospect theory investor overweights small probabilities, he loves lottery-like wealth distributions, and is therefore willing to pay a very high price for the skewed security. Of course, since the equilibrium we construct assumes the existence of short-sale constraints, just how overpriced the skewed security is, depends on how high short-selling costs are.

Our result that a skewed security can become overpriced is by no means an obvious one. An investor who overweights low probabilities will, of course, value a skewed *portfolio* highly; but what is surprising is that he also values a skewed *security* highly, even if that security is small and independent of other risks. Indeed, as we discuss below, such a security would *not* earn a low average return in a traditional expected utility model of skewness preference.

It is natural to ask whether, for the parameter pair $(\sigma_M, L) = (0.15, 10)$, there are any equilibria other than the heterogeneous holdings equilibrium described above. While it is difficult to give a definitive answer, we can at least show that, for these parameters, there is no representative agent equilibrium. In such an equilibrium, all investors would need to be happy to hold an infinitesimal amount of the new security, which, in turn, means that the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ would need to have a derivative of 0 at $x = 0$. To achieve this, we would need to set $q = 0.0926$, which corresponds to the dashed line in Figure 3. This line also shows, however, that $x = 0$ is then no longer a global optimum: all investors would prefer a substantial positive position in the new security, making it impossible to clear the market.[7]

5.2 How does expected return vary with skewness?

By trying many different values of L, we have found that our results for $L = 10$ are typical of the results for *any* large L: a heterogeneous holdings equilibrium can be constructed, but a representative agent equilibrium cannot. For low values of L, however, we have found that the opposite is true: a representative agent equilibrium can be constructed, but a heterogeneous holdings equilibrium cannot. The reason the heterogeneous holdings equilibrium breaks down is that if the new security is not sufficiently skewed, no position in it, however large, adds enough skewness to the investor's portfolio to compensate for the lack of diversification the position entails.

For example, suppose that, as before, $\sigma_M = 0.15$, so that, once again, $\mu_M = 0.0747$, but that instead of $L = 10$, we set $L = 4$. Figure 4 shows the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ for various values of q, namely $q = 0.15$ (dashed line), $q = 0.206$ (solid line), and $q = 0.35$ (dash-dot line). While these lines correspond to only three values of q, they hint at the outcome of a more extensive search, namely that no value of q can deliver two global optima.

For this example, we can only obtain a representative agent equilibrium. This corresponds to the solid line in Figure 4, for which the goal function has a derivative of 0 at $x = 0$, so that all investors are happy to hold an infinitesimal amount of the skewed security. For this value of q, the expected excess return of the new security is, from equation (28), equal to

$$E(\widehat{R}_n) = (0.206)(4) - (0.794)(1.02) = 0.$$

The fact that the new security earns an expected excess return exactly equal to zero is no coincidence. The following proposition shows that whenever a representative agent

[7]Given the scale of Figure 3, it is hard to tell whether the dashed line really does have a derivative of 0 at $x = 0$. Magnifying the left side of the graph confirms that the derivative *is* 0 at $x = 0$, although it quickly turns negative as x increases.

equilibrium exists, the expected excess return on the new security is *always* equal to zero. In other words, in a representative agent setting, cumulative prospect theory assigns the new security the *same* average return that a concave expected utility specification would.

Proposition 3: *Consider an agent who holds a portfolio with return $\tilde{R} \equiv \hat{R} + R_f$. Suppose that he adds a small amount of the new security to his current holdings and finances this by borrowing, so that his portfolio return becomes $\hat{R} + \epsilon \hat{R}_n$. If $\hat{R}$ and $\hat{R}_n$ are independent, and* if *$\hat{R}$ has a distribution density function with $\sigma(\hat{R}) > 0$, then*

$$\lim_{\epsilon \to 0} \frac{V(\hat{R} + \epsilon \hat{R}_n) - V(\hat{R})}{\epsilon} = E(\hat{R}_n) V'(\hat{R}), \tag{29}$$

where, with some abuse of notation, $V'(\cdot)$ is defined as

$$\begin{aligned} & V'(\hat{R}) \\ \equiv \ & \lim_{x \to 0} \frac{V(\hat{R} + x) - V(\hat{R})}{x} \\ = \ & \int_{-\infty}^{0} w'(P(R)) P'(R) dv(R) + \int_{0}^{\infty} w'(1 - P(R)) P'(R) dv(R) > 0, \end{aligned} \tag{30}$$

*where $P(R) \equiv$*Prob*$(\hat{R} \leq R)$.*

Proof of Proposition 3: See the Appendix.

In any representative agent equilibrium, we need $V(\widehat{R} + x\widehat{R}_n)$ to have a local optimum at $x = 0$. From the proposition, this means that $E(\widehat{R}_n) = 0$.

In summary, we have found that for high values of L, equilibrium involves heterogeneous portfolio holdings and a negative expected excess return for the skewed security, while for low values of L, investors hold identical portfolios and the expected excess return is zero. Figure 5 shows the quantitative relationship between expected return and L. For high values of L, we obtain a heterogeneous holdings equilibrium and expected return varies inversely with L: when L is high, the new security can add a large amount of skewness to the investor's portfolio. It is therefore more valuable to him and he lowers the expected return he requires on it. Once the skewness of the new security falls below a certain level – in our case, once it falls below $L = 8$ – a heterogenous holdings equilibrium can no longer be constructed. There is, however, a representative agent equilibrium, and here, as seen in Proposition 3, the new security earns a zero expected excess return.

5.3 Relation to the earlier literature on skewness

Through the probability weighting function, cumulative prospect theory investors exhibit a preference for skewness. There are already a number of papers that analyze the implications of skewness-loving preferences (Kraus and Litzenberger, 1976). We note, however, that the pricing effect we have just demonstrated is new to the skewness literature. Earlier papers have shown that, in the presence of skewed securities, coskewness with the market can earn a risk premium. We show that skewness itself, and not just coskewness with the market, can earn a premium. For example, in our economy, the skewed security earns a negative average excess return even though it is small and its returns are independent of other risks; in other words, even though its coskewness with the market is zero.

How do we obtain this new effect? The earlier skewness literature only considers economies in which all investors hold diversified portfolios. In such economies, only coskewness with the market portfolio earns a premium. We consider a more general economy with heterogeneous portfolio holdings in which some investors hold undiversified positions. In such an economy, it is not only coskewness with the market that matters, but also a security's skewness itself.

The analysis above also shows, however, that under CPT preferences, the relationship between skewness and expected return is highly nonlinear: only securities with a high degree of skewness earn a negative expected excess return, while those with merely moderate skewness have an expected excess return of zero.

5.4 How does expected return vary with the preference parameters?

So far, we have looked at how the expected return of the skewed security depends on L, while keeping the preference parameters fixed at the values estimated by Tversky and Kahneman (1992), namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$. We now fix L, and examine the effect of varying α, δ, and λ. The three panels in Figure 6 plot the expected excess return of the new security against each of these parameters in turn, holding the other two constant.

The top left panel shows that, as δ increases, the expected return of the new security also rises. A low value of δ means that the investor weights small probabilities particularly heavily and therefore that he is strongly interested in a positively skewed portfolio. Since the new security offers a way of constructing such a portfolio, it is very valuable, and the investor is willing to hold it in exchange for a very low average return. Once δ rises above 0.68, however, no heterogeneous holdings equilibrium is possible: by this point, the investor does not care enough about having a positively skewed portfolio for him to want to take on an undiversified position in the new security. Only a representative agent equilibrium is possible, and in such an equilibrium, the new security earns a zero average excess rate of return.

The graph at top right shows that, as λ increases, the expected return on the new security also rises. The parameter λ governs the investor's aversion to losses. In order to add skewness to his portfolio, the investor needs to hold a large, undiversified position in the new security. As λ increases, he becomes less willing to tolerate the high volatility of this undiversified position and is therefore only willing to hold the skewed security in return for a higher expected return. Once λ rises above 2.4, no position in the new security, however large, contributes sufficient skewness to offset the painful swings in the overall portfolio. Beyond this point, only a representative agent equilibrium is possible.

Finally, the graph in the lower left panel shows that, as α falls, the expected return on the new security goes up. A lower α means that the value function in the region of gains, depicted in Figure 1, is more concave. This means that the investor derives less utility from "hitting the jackpot" and therefore, less utility from having a positively skewed portfolio. As a result, the new security is less useful to him, and he is only willing to hold it in exchange for a higher average return.

5.5 Multiple skewed securities

We now show that our basic result – that a positively skewed security can earn negative average excess returns, even when small and independent of other risks – continues to hold even when there are several skewed securities in the economy. Suppose, for example, that there are $M > 1$ identical, positively skewed securities, each of them satisfying Assumptions 13 and 14, and each of them distributed according to (26). We find that, so long as a heterogeneous holdings equilibrium exists when there is just one skewed security, then such an equilibrium also exists when there are M of them. Moreover, the equilibrium takes a very simple form: some investors hold a large, undiversified position in just the first skewed security, others hold a large undiversified position in just the second skewed security, and so on for each of the skewed securities; and the remaining investors hold no position at all in any of the skewed securities. As before, the heterogeneous holdings are the result of non-unique optimal portfolios: holding no position at all in any of the skewed securities and holding a substantial position in any one of them, are both global optima. And, as before, each of the positively skewed securities earns a negative risk premium in equilibrium.

To be sure that such a heterogeneous holdings equilibrium exists, we need to check that, in the proposed equilibrium, investors do indeed prefer a substantial position in just one skewed security, rather than a diversified position in several of them. We demonstrate this for $M = 2$. Suppose that, as in our example in Section 5.1, $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$ and $\sigma_M = 0.15$, so that $\mu_M = 0.0747$. Now suppose that there are two positively skewed securities, each of which is small and independent of all other securities and each of which has the distribution in (26) with $L = 10$. We denote the excess return of the first skewed security as $\widehat{R}_{n,1}$ and that of the second as $\widehat{R}_{n,2}$.

For $q = 0.087$ – in other words, for the value of q that gives a heterogeneous holdings equilibrium in the case of one skewed security – Figure 7 plots the goal function $V(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2})$, the utility of adding a position x_1 in security 1 and x_2 in security 2 to the market portfolio. To compute it, note that

$$V(\hat{R}_M + x_1\hat{R}_{n,1} + x_2\widehat{R}_{n,2}) = -\int_{-\infty}^{0} w(P_{x_1,x_2}(R))dv(R) + \int_{0}^{\infty} w(1 - P_{x_1,x_2}(R))dv(R) \quad (31)$$

where

$$\begin{aligned}
P_{x_1,x_2}(R) &\equiv \Pr(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2} \leq R) \\
&= \Pr(\widehat{R}_{n,1} = L)\Pr(\widehat{R}_{n,2} = L)\Pr(\widehat{R}_M \leq R - xL - xL) \\
&+ \Pr(\widehat{R}_{n,1} = L)\Pr(\widehat{R}_{n,2} = -R_f)\Pr(\widehat{R}_M \leq R - xL + xR_f) \\
&+ \Pr(\widehat{R}_{n,1} = -R_f)\Pr(\widehat{R}_{n,2} = L)\Pr(\widehat{R}_M \leq R + xR_f - xL) \\
&+ \Pr(\widehat{R}_{n,1} = -R_f)\Pr(\widehat{R}_{n,2} = -R_f)\Pr(\widehat{R}_M \leq R + xR_f + xR_f) \\
&= q^2 N(\frac{R - 2xL - \mu_M}{\sigma_M}) + 2q(1-q)N(\frac{R + xR_f - xL - \mu_M}{\sigma_M}) \qquad (32) \\
&\quad +(1-q)^2 N(\frac{R + 2xR_f - \mu_M}{\sigma_M}).
\end{aligned}$$

The solid line in Figure 7 corresponds to $x_2 = 0$, and is therefore identical to the solid line in Figure 3. The remaining lines correspond to $x_2 = 0.065$ (dashed line), $x_2 = 0.13$ (dash-dot line) and $x_2 = 0.195$ (dotted line).

Figure 7 shows that $q = 0.087$ delivers a heterogeneous holdings equilibrium, not only for the case of one skewed security, but also for the case of two. Since $V(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2})$ is negative whenever x_1 and x_2 are both positive, the investor prefers a large undiversified position in just one skewed security to a diversified position in two of them. In other words, investors who allocate to the skewed securities at all, allocate to just one of them. Each skewed security is therefore again overpriced and, as before, earns a negative expected excess return of -7.23 percent.

The intuition behind Figure 7 is that, by diversifying across several skewed securities, the investor lowers the volatility of his overall portfolio – which is good – but also lowers its skewness – which is bad. Which of these two forces dominates depends on how the security returns are distributed. Our analysis shows that, for the binomial distribution in (26), skewness falls faster than volatility, making diversification unattractive.

6 Applications

6.1 The pricing of IPOs

The overpricing of positively skewed securities seen in Section 5 suggests a new way of thinking about a number of puzzling asset market phenomena, most notably the low average returns on IPOs (Ritter, 1991). IPOs have a positively skewed return distribution, probably because, being young firms, a large fraction of their value is in the form of growth options. Our analysis then shows that, in an economy populated by investors with cumulative prospect theory preferences, IPO securities may become overpriced and earn low average returns. Under cumulative prospect theory, then, the historical performance of IPOs may not be so puzzling.

To be sure that our theoretical results can indeed be applied to IPOs, we need to check that IPO returns have *enough* skewness to justify a low average return. After all, our earlier analysis – Figure 5, in particular – shows that investors with the preferences in (11)-(13) overprice highly skewed securities, but not those with merely moderate skewness. As a result, we need to check that there is enough skewness in IPO returns to make a heterogeneous holdings equilibrium possible.

To do this, we replace the binomial distribution assumed in Section 5.1 with the actual empirical distribution of IPO returns, estimated from 1525 IPOs between 1975 and 1984, Ritter's (1991) original sample. Since the underperformance of IPOs is a feature of long-horizon returns, we work with 3-year buy-and-hold returns. These returns are plotted in Figure 8.

In principle, we could replace the binomial distribution (26) with the discrete distribution that assigns each of the 1525 returns in our sample equal probability:

$$(\frac{1}{1525}, \widehat{R}_{n,1525}^{(1)}; \frac{1}{1525}, \widehat{R}_{n,1525}^{(2)}; \ldots; \frac{1}{1525}, \widehat{R}_{n,1525}^{(1525)}), \tag{33}$$

where $\{\widehat{R}_{n,1525}^{(j)}\}_{j=1,\ldots,1525}$ are the excess buy-and-hold returns, arranged in ascending order, so that $\widehat{R}_{n,1525}^{(1)} \leq \widehat{R}_{n,1525}^{(2)} \leq \ldots \leq \widehat{R}_{n,1525}^{(1525)}$.

For the sake of computational tractability, we replace the binomial distribution (26) with a $1525/5 = 305$-point discrete distribution that "aggregates" the 1525-point distribution 5 points at a time:

$$(\frac{1}{305}, \widehat{R}_{n,305}^{(1)}; \frac{1}{305}, \widehat{R}_{n,305}^{(2)}; \ldots; \frac{1}{305}, \widehat{R}_{n,305}^{(305)}), \tag{34}$$

where

$$\widehat{R}_{n,305}^{(j)} = \frac{1}{5}\sum_{l=0}^{4} \widehat{R}_{n,1525}^{(5j-l)}.$$

As before, we take the preference parameters to be $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$. Since we are working with 3-year IPO returns, we reset all security return parameters to a 3-year time frame. In particular, we set σ_M to 0.26, rather than to 0.15, so as to match the standard deviation of 3-year market returns. Solving equation (26) gives an average 3-year excess market return of $\mu_M = 12.96$ percent.

We now search for a constant c, such that conditions (21)-(22) are satisfied when $\widehat{R}_n$ has the probability distribution

$$(\frac{1}{305}, \widehat{R}_{n,c}^{(1)}; \frac{1}{305}, \widehat{R}_{n,c}^{(2)}; \ldots; \frac{1}{305}, \widehat{R}_{n,c}^{(305)}), \tag{35}$$

where

$$\widehat{R}_{n,c}^{(j)} = \max(\widehat{R}_{n,305}^{(j)} - c, -R_f).$$

In other words, we look for a constant c, so that once the distribution of IPO returns has been shifted by c, a heterogeneous holdings equilibrium obtains. We define $\widehat{R}_{n,c}^{(j)}$ as $\max(\widehat{R}_{n,305}^{(j)} - c, -R_f)$ rather than simply as $\widehat{R}_{n,305}^{(j)} - c$, so as to ensure truncation at $-R_f$, the lowest value $\widehat{R}_n$ can take.

In order to compute the goal function $V(\widehat{R}_M + x\widehat{R}_n)$, we need $P_x(R)$, which is given by

$$\begin{aligned} P_x(R) &\equiv P(\widehat{R}_M + x\widehat{R}_n \leq R) \\ &= \sum_{j=1}^{305} \Pr(\widehat{R}_n = \widehat{R}_{n,c}^{(j)}) \Pr(\widehat{R}_M \leq R - x\widehat{R}_{n,c}^{(j)}) \\ &= \sum_{j=1}^{305} \frac{1}{305} N(\frac{R - x\widehat{R}_{n,c}^{(j)} - \mu_M}{\sigma_M}). \end{aligned} \tag{36}$$

Finally, we set the 3-year risk-free interest rate to $R_f = 1.02^3 = 1.0612$.

If our cumulative prospect theory framework predicted an average return for IPOs exactly equal to the historical value, c would equal 0. We should not expect this to be the case, however: in part because probability weighting is unlikely to be the only force driving the average IPO return; in part because even if it were, c would differ from 0 simply because of sampling error; and in part because the mean excess market return μ_M predicted by our model differs from the actual market return over the 1975-1984 period.

We find that IPO returns *do* have sufficient skewness to allow for a heterogeneous holdings equilibrium. Figure 9 shows the goal function for this equilibrium, which corresponds to $c = -0.434$. As before, some investors hold no position at all in the IPO security; others hold a large, undiversified position in it, leading it to be overpriced and to earn a negative average excess return, $E(\widehat{R}_n)$, which, in this equilibrium, equals -0.0843.

We also confirm, using a calculation analogous to that of Section 5.5, that this result is robust to the introduction of many IPO securities into the economy. Even when there are several such securities, each of them independent of other risks and each of them distributed as in (35), the investor prefers a substantial position in one IPO to a diversified position in several of them.

Our analysis shows that a positively skewed security with *no* systematic risk earns an excess return of -8.43 percent, 8.43 percent lower than the 0 percent excess average return it would earn in a representative agent, expected utility world. In a more realistic model, the IPO security would have systematic risk of its own – enough systematic risk to earn the market return, say, in an expected utility framework. To the extent that the results from our simple model translate to this richer environment, our analysis suggests that in this latter economy, the preferences in (11)-(13) would assign the IPO security an average return several percentage points below the market return. Cumulative prospect theory can therefore go some way toward explaining the historical performance of IPOs.

The potential application of cumulative prospect theory to the overpricing of IPOs has also been noted by Brav and Heaton (1996) and Ma and Shen (2003). However, in writing down the equilibrium conditions for their quantitative work, these papers assume that investors engage in "narrow framing" – in other words, that investors get utility *directly* from individual stocks that they own, even if these stocks are only a small part of their portfolios. Ma and Shen (2003), for example, argue that, in equilibrium, the prospective utility of a typical IPO security should equal the prospective utility of a typical seasoned stock. For an investor to care about the utility of a single security in this way only makes sense in the presence of narrow framing.

While narrow framing may well occur in practice, we believe that, as a first step, it is better to analyze the implications of probability weighting under the traditional assumption that investors get utility only from overall *portfolio* returns, not from individual security returns. Our equilibrium conditions (21)-(22) reflect this assumption.

6.2 Other applications

Our results may also shed light on the recently-documented examples of "equity stubs" with remarkably low valuations (Mitchell, Pulvino and Stafford, 2002, Lamont and Thaler, 2003). These are cases of firms with publicly traded subsidiaries in which the subsidiaries make up a surprisingly large fraction of the value of the parent company; in extreme cases, more than 100 percent of the value of the parent company, so that the subsidiary is worth more than the parent, and so that the equity stub – the claim to the parent company's businesses outside of the subsidiary – has a negative value.

Our analysis cannot explain negative stub values, but it may explain stub values that are surprisingly low, albeit positive. If a subsidiary is valued mainly for its growth options, its returns may be positively skewed, leading investors to overprice it relative to its parent, and thereby generating a low stub value. Consistent with this, in most of the examples listed by Mitchell, Pulvino and Stafford (2002), the subsidiary's business activities involve newer technologies – and therefore, in all likelihood, more growth options – than do the parent company's.

Of course, since the subsidiary forms part of the parent company, its growth options may also give the parent company a positively skewed distribution. Even if this is the case, however, the parent company will still be much less skewed than the subsidiary and, as seen in Section 5.2, it is only high levels of skewness that are overpriced; more moderate levels of skewness are fairly valued. The subsidiary can therefore be overpriced even relative to its parent.

Our analysis also predicts that a firm can create value by spinning off subsidiaries that are valued mainly for their growth options. If a subsidiary of this kind is traded as a separate entity, its stock is likely to have a positively skewed return distribution. If investors overprice such securities, the firm will be more valuable when its parts are traded separately, than when they are traded as a single bundle.

7 Conclusion

As part of their cumulative prospect theory, Tversky and Kahneman (1992) argue that, when people evaluate risk, they transform objective probabilities via a weighting function which, among other features, overweights small probabilities. We investigate the implications of cumulative prospect theory for the pricing of financial securities, paying particular attention to the effects of the weighting function. Under cumulative prospect theory, the CAPM can hold when securities are normally distributed; but a positively skewed security can become overpriced and earn very low average returns, even if small and independent of other risks, and even if just one of many skewed securities in the economy. We apply the last result to the pricing of IPOs and to the valuation of equity stubs. Using data on the skewness of IPO returns, we show that investors with cumulative prospect theory preferences calibrated to experimental evidence would require an average return on IPOs that is several percentage points below the market return. Under cumulative prospect theory, then, the historical underperformance of IPOs may not be so puzzling.

8 Appendix

Proof of Lemma 1: Since $w(1-P(\cdot))$ is right-continuous and $v(\cdot)$ is continuous, we can integrate (12) by parts to obtain

$$V(\hat{W}^+) = [-v(x)w(1-P(x))]_{x=0}^{x=\infty} + \int_0^\infty w(1-P(x))dv(x). \tag{37}$$

Using Chebychev's inequality,

$$P\left[|\hat{W} - E(\hat{W})| \geq Z\right] \leq \frac{\text{Var}(\hat{W})}{Z^2}, \tag{38}$$

we have

$$1 - P(x) \leq P\left[|\hat{W} - E(\hat{W})| \geq x + E(\hat{W})\right] \leq \frac{\text{Var}(\hat{W})}{(x + E(\hat{W}))^2}. \tag{39}$$

Combining this with Assumptions 1-5, we can show that there exists $A > 0$ such that

$$v(x)w(1-P(x)) \leq Ax^{\alpha-2\delta} \to 0, \quad \text{as } x \to \infty. \tag{40}$$

The first term on the right-hand side of equation (37) is therefore zero, and so equation (16) is valid. A similar argument can be used to establish equation (17).

Before proving Proposition 1, we first prove the following useful lemma.

Lemma 2: *Consider the preferences in (11)–(13) and suppose that Assumptions 2-5 hold. If $\hat{W} \sim N(\mu_W, \sigma_W^2)$ is normally distributed, then $V(\hat{W}) = F(\mu_W, \sigma_W^2)$. Moreover, $F(\mu_W, \sigma_W^2)$ is strictly increasing in μ_W for any σ_W^2.*

Proof of Lemma 2: Since every normal distribution is fully specified by its mean and variance, we can write $V(\hat{W}) = F(\mu_W, \sigma_W^2)$. Proposition 2 implies that $F(\mu_W, \sigma_W^2)$ is strictly increasing in μ_W.

Proof of Proposition 1: We ignore the violation of limited liability and assume that all securities have positive prices in equilibrium. Consider the mean/standard deviation plane. For any set of positive prices for the J securities, Assumption 8 implies that we can construct a hyperbola representing the mean-variance (MV) frontier for the J risky assets. When a risk-free security is available, the MV frontier is the tangency line from the risk-free return to the hyperbola, plus the reflection of this tangency line off the vertical axis. The MV *efficient* frontier is the upper of these two lines. The return of the tangency portfolio, composed only of the J risky securities, is denoted $\tilde{R}_T$.

By Lemma 2, each agent's portfolio is on the MV efficient frontier, and so can be characterized as $\tilde{R} = R_f + \theta(\tilde{R}_T - R_f)$, with θ denoting the portfolio weight in the tangency portfolio. Since all agents – even if they have different parameters in their preferences – choose the same sign for θ, market clearing implies that $\theta > 0$, so $\tilde{R}_T$ has to be on the upper half of the hyperbola in equilibrium. We consider this case only. In this case, $E(\hat{R}_T) > 0$, since R_f has to be lower than the expected return of the minimum-variance portfolio, namely the left-most point of the hyperbola.

With $\hat{R}_T \equiv \tilde{R}_T - R_f$, each agent's time-1 wealth is

$$\begin{aligned} \tilde{W} &= W_0(R_f + \theta(\tilde{R}_T - R_f)) & (41) \\ \hat{W} &= W_0\theta\hat{R}_T, & (42) \end{aligned}$$

and for $\theta > 0$, utility, which is homogeneous to degree α, is given by

$$U(\theta) = W_0^\alpha \theta^\alpha V(\hat{R}_T) \tag{43}$$

where

$$V(\hat{R}_T) = -\int_{-\infty}^{0} w(P(\hat{R}_T))dv(\hat{R}_T) + \int_{0}^{\infty} w(1 - P(\hat{R}_T))dv(\hat{R}_T). \tag{44}$$

The optimal solution for the agent is therefore

$$\theta = \begin{cases} 0, & \text{for } V(\hat{R}_T) < 0 \\ \text{any } \theta \geq 0, & \text{for } V(\hat{R}_T) = 0 \\ \infty, & \text{for } V(\hat{R}_T) > 0. \end{cases} \tag{45}$$

In equilibrium, the aggregate demand portfolio, given by (45), is the same as the aggregate supply portfolio, namely the market portfolio. So $\theta = 1$ in equilibrium, and $\tilde{R}_T = \tilde{R}_M$. We therefore have $V(\hat{R}_M) = 0$ and $E(\hat{R}_M) > 0$ in equilibrium. Finally, $\tilde{R}_T = \tilde{R}_M$ implies equations (18)-(20) for all securities. For example, this can be shown by noting that portfolio $\tilde{R}_M + x(\tilde{R}_i - R_f)$ attains its highest Sharpe ratio at $x = 0$.

Proof of Proposition 3: The Gateaux derivative in (30) follows from

$$\frac{\partial V(\hat{R} + x)}{\partial x}|_{x=0} = \frac{\partial}{\partial x}|_{x=0}\left[-\int_{-\infty}^{0} w(P(R - x))dv(R) + \int_{0}^{\infty} w(1 - P(R - x))dv(R)\right]. \tag{46}$$

To show the main result, we have, to the first order of ϵ,

$$\begin{aligned} \delta V &\equiv V(\hat{R} + \epsilon\hat{R}_n) - V(\hat{R}) \\ &\approx -\int_{-\infty}^{0} w'(P(R))\delta P(R)dv(R) - \int_{0}^{\infty} w'(1 - P(R))\delta P(R)dv(R), & (47) \end{aligned}$$

where, again to the first order of ϵ,

$$\delta P(R) \equiv P(\hat{R} + \epsilon\hat{R}_n \leq R) - P(\hat{R} \leq R)$$

$$
\begin{aligned}
&= E\left[1(\hat{R}+\epsilon\hat{R}_n \le R) - 1(\hat{R}\le R)\right] \\
&= E\left[-1(\epsilon\hat{R}_n > 0)1(R-\epsilon\hat{R}_n < \hat{R} \le R) + 1(\epsilon\hat{R}_n < 0)1(R < \hat{R} \le R-\epsilon\hat{R}_n)\right] \\
&\approx E\left[-1(\epsilon\hat{R}_n > 0)f(R|\hat{R}_n)(\epsilon\hat{R}_n) + 1(\epsilon\hat{R}_n < 0)f(R|\hat{R}_n)(-\epsilon\hat{R}_n)\right] \\
&= E\left[f(R|\hat{R}_n)(-\epsilon\hat{R}_n)\right] \\
&= -\epsilon E(\hat{R}_n)f(R), \qquad (48)
\end{aligned}
$$

where $f(R|\hat{R}_n)$ and $f(R)$ are the conditional and unconditional probability densities of $\hat{R}$ at R, respectively, and where the last equality in equation (48) follows from the independence assumption. Putting equation (48) into equation (47), and applying equation (30), we obtain equation (29).

9 References

Ang, A., Bekaert G. and J. Liu (2002), "Why Stocks May Disappoint," Working paper, Columbia University.

Barberis, N., Huang M., and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* 116, 1-53.

Barberis, N., Huang M., and R. Thaler (2003), "Individual Preferences, Monetary Gambles and the Equity Premium," Working paper, University of Chicago.

Benartzi, S. and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 75-92.

Brav, A. and J.B. Heaton (1996), "Explaining the Underperformance of IPOs: A Cumulative Prospect Theory Approach," Working paper, Duke University.

Dahlquist M. and F. De Jong (2004), "Pseudo Market Timing: Fact or Fiction?," Working paper, Stockholm School of Economics.

Epstein, L. and S. Zin (1990), "First-order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics* 26, 387-407.

Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263-291.

Kraus, A. and R. Litzenberger (1976), "Skewness Preference and the Valuation of Risk Assets," *Journal of Finance* 31, 1085-1100.

Lamont O., and R. Thaler (2003), "Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs," *Journal of Political Economy* 111, 227-268.

Levy, H., De Giorgi E., and T. Hens (2003), "Prospect Theory and the CAPM: A Contradiction or Co-existence?" Working paper, University of Zurich.

Ma T., and Y. Shen (2003), "Prospect Theory and the Long-run Performance of IPO Stocks," Working paper, University of Utah.

Mitchell M., Pulvino P., and E. Stafford (2002), "Limited Arbitrage in Equity Markets," *Journal of Finance* 57, 551-584.

Rabin, M. (2000), "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica* 68, 1281-1292.

Ritter, J. (1991), "The Long-run Performance of Initial Public Offerings," *Journal of Finance* 46, 3-27.

Schultz, P. (2003), "Pseudo Market Timing and the Long Run Performance of IPOs," Journal of Finance 58, 483-518.

Tversky, A., and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297-323.

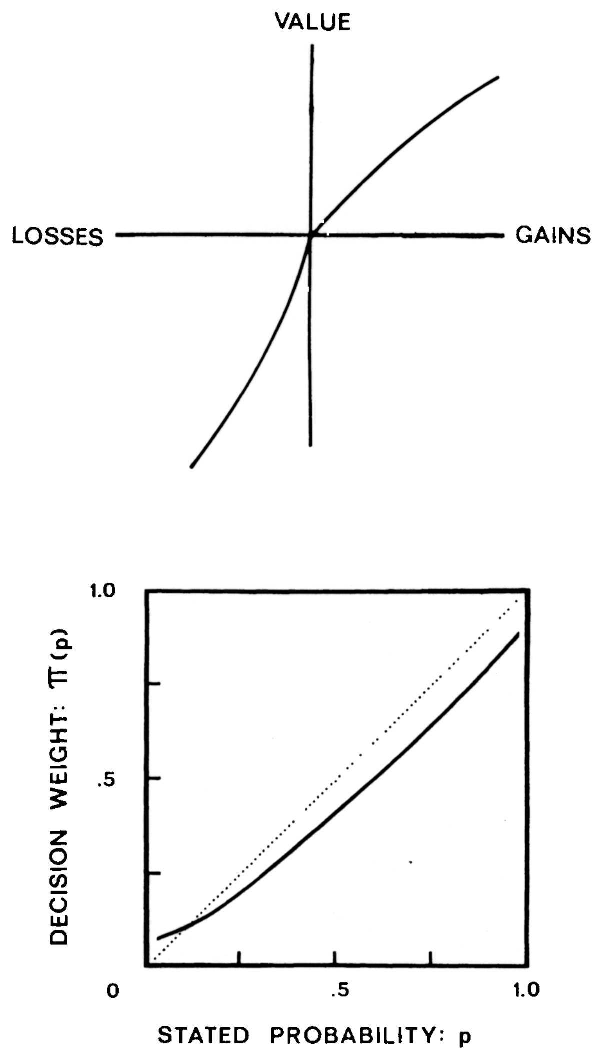

Figure 1. The two panels show Kahneman and Tversky's (1979) proposed value function $v(\cdot)$ and probability weighting function $\pi(\cdot)$.

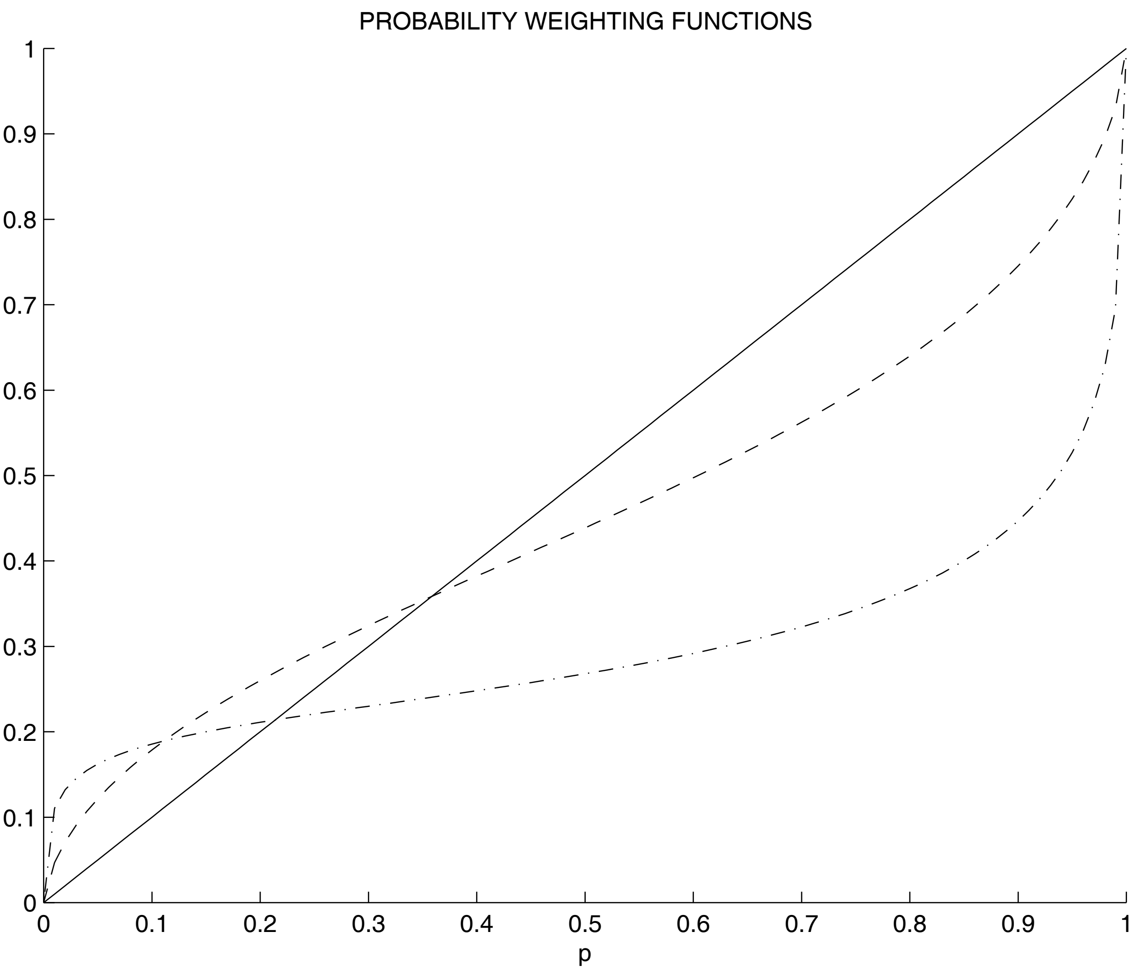

Figure 2. The figure shows the form of the probability weighting function proposed by Tversky and Kahneman (1992), for parameter values $\delta = 0.65$ (dashed line), $\delta = 0.4$ (dash-dot line), and $\delta = 1$, which corresponds to no probability weighting at all (solid line). p is the objective probability.

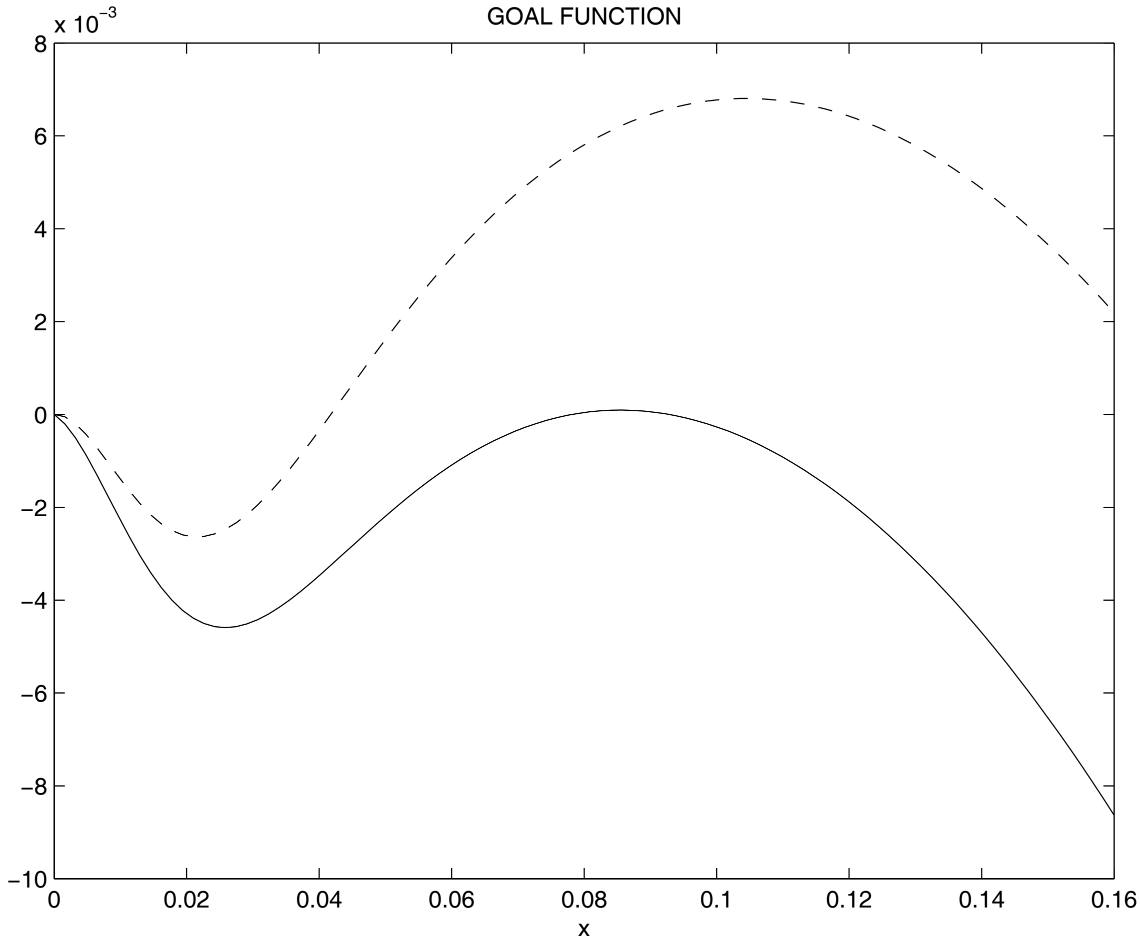

Figure 3. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x in a positively-skewed security to his current holdings of a normally distributed market portfolio. The dashed line corresponds to a higher mean return on the skewed security.

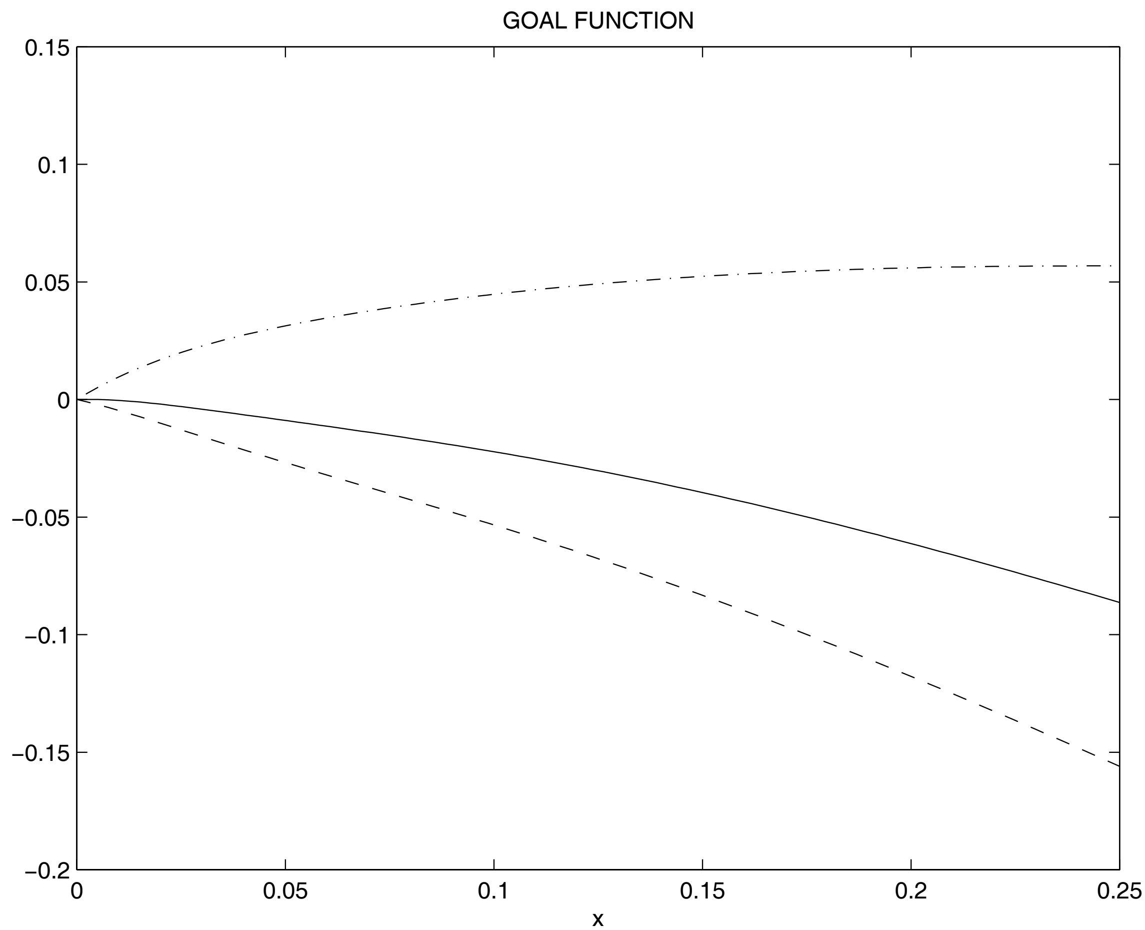

Figure 4. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x in a positively-skewed security to his current holdings of a normally distributed market portfolio. The three lines correspond to different mean returns on the skewed security.

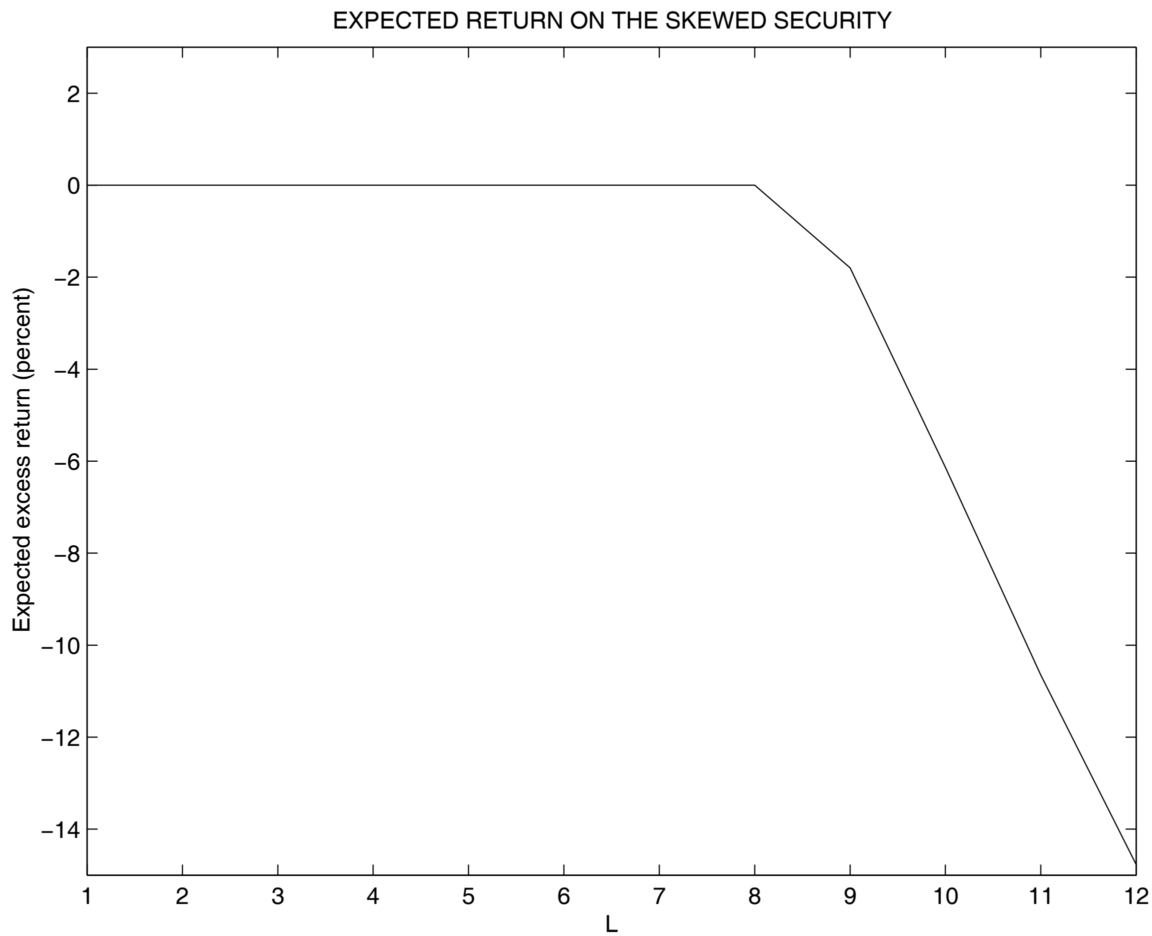

Figure 5. The figure shows the expected return in excess of the risk-free rate earned by a small, independent, positively skewed security in an economy populated by investors who judge gambles according to cumulative prospect theory, plotted against a parameter of the the security's return distribution. The security earns a gross return of 0 with high probability and of L with low probability.

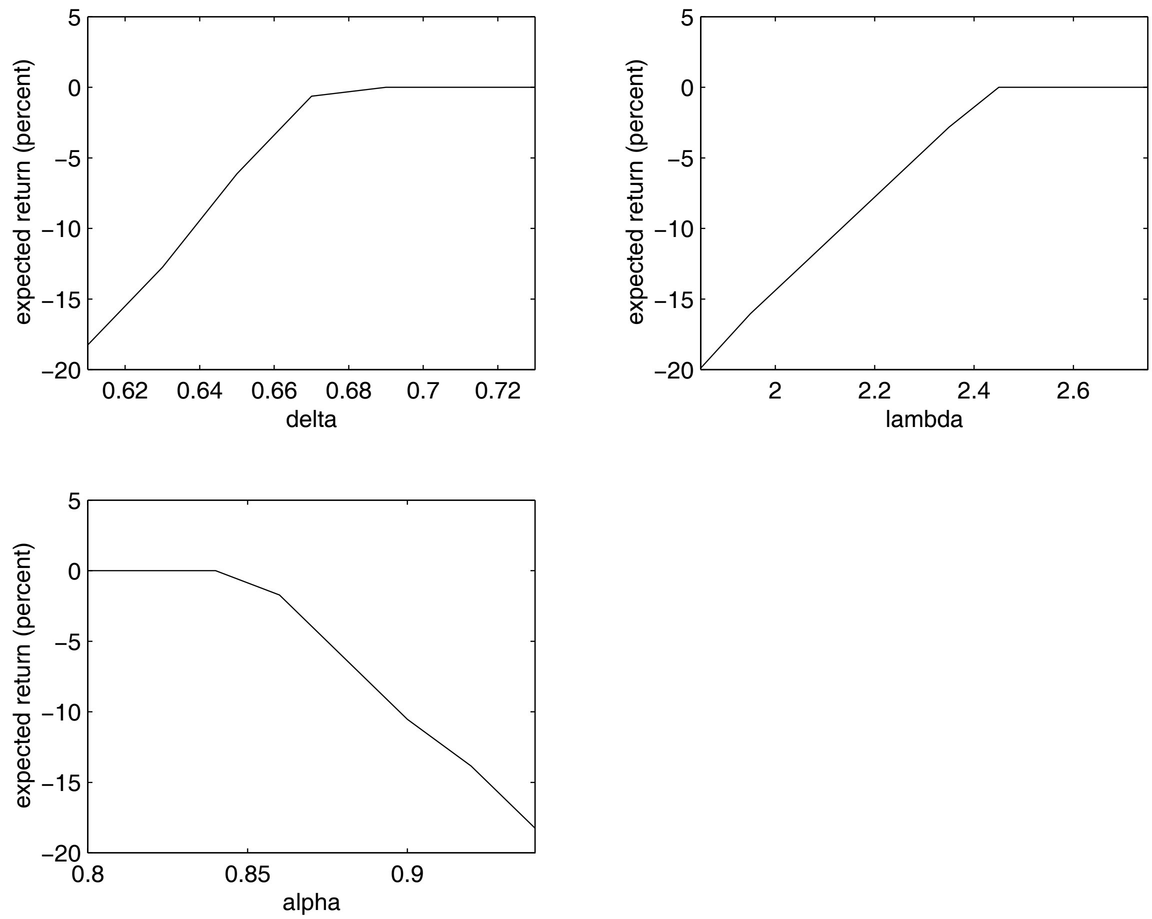

Figure 6. The figure shows the expected return in excess of the risk-free rate earned by a small, independent, positively skewed security in an economy populated by investors who judge gambles according to cumulative prospect theory, plotted against parameters of investors' utility functions. As δ falls, investors overweight small probabilities more heavily; as λ increases, they become more sensitive to losses; and as α falls, their marginal utility from gains falls.

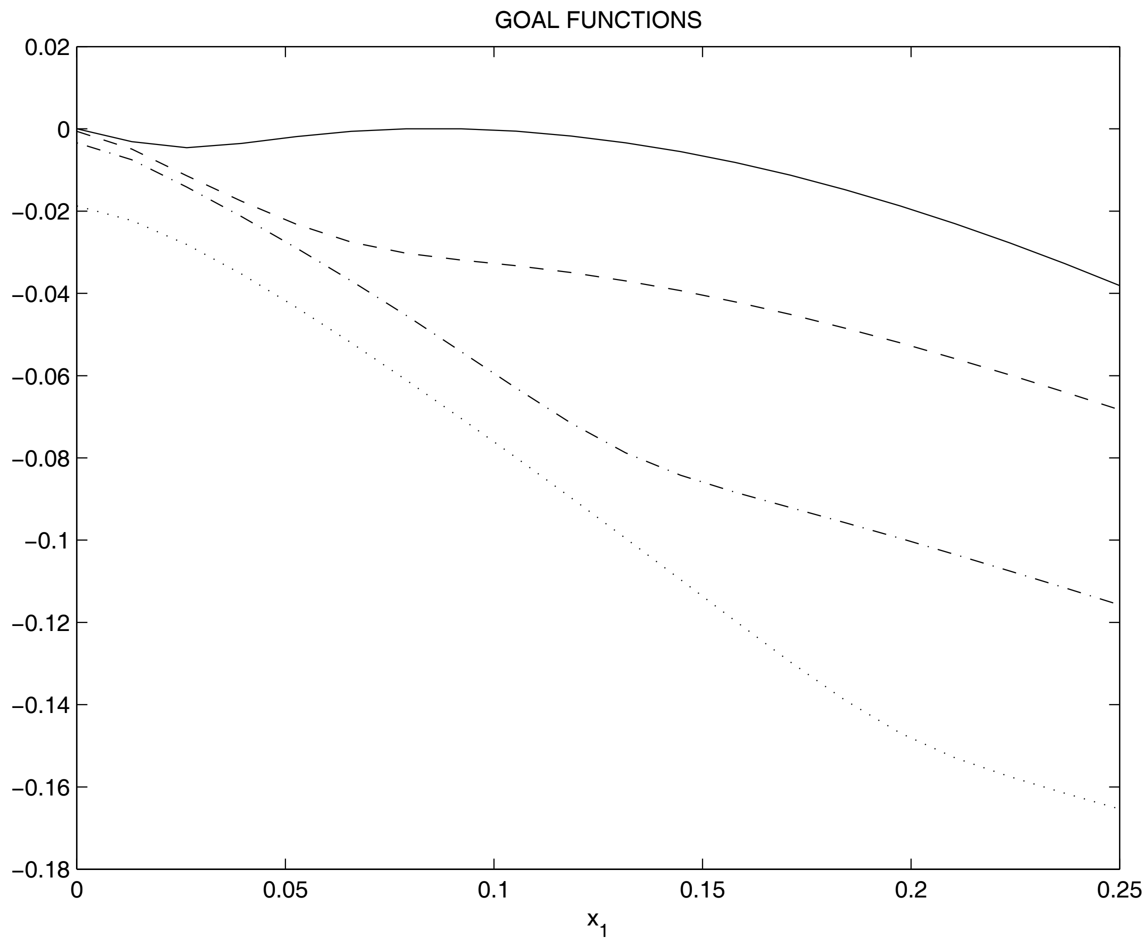

Figure 7. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x_1 in a positively-skewed security and a position x_2 in another positively-skewed security to his current holdings of a normally distributed market portfolio. The four lines correspond to different values of x_2: $x_2 = 0$ (the solid line), $x_2 = 0.065$ (the dashed line), $x_2 = 0.13$ (the dash-dot line), $x_2 = 0.195$ (the dotted line).

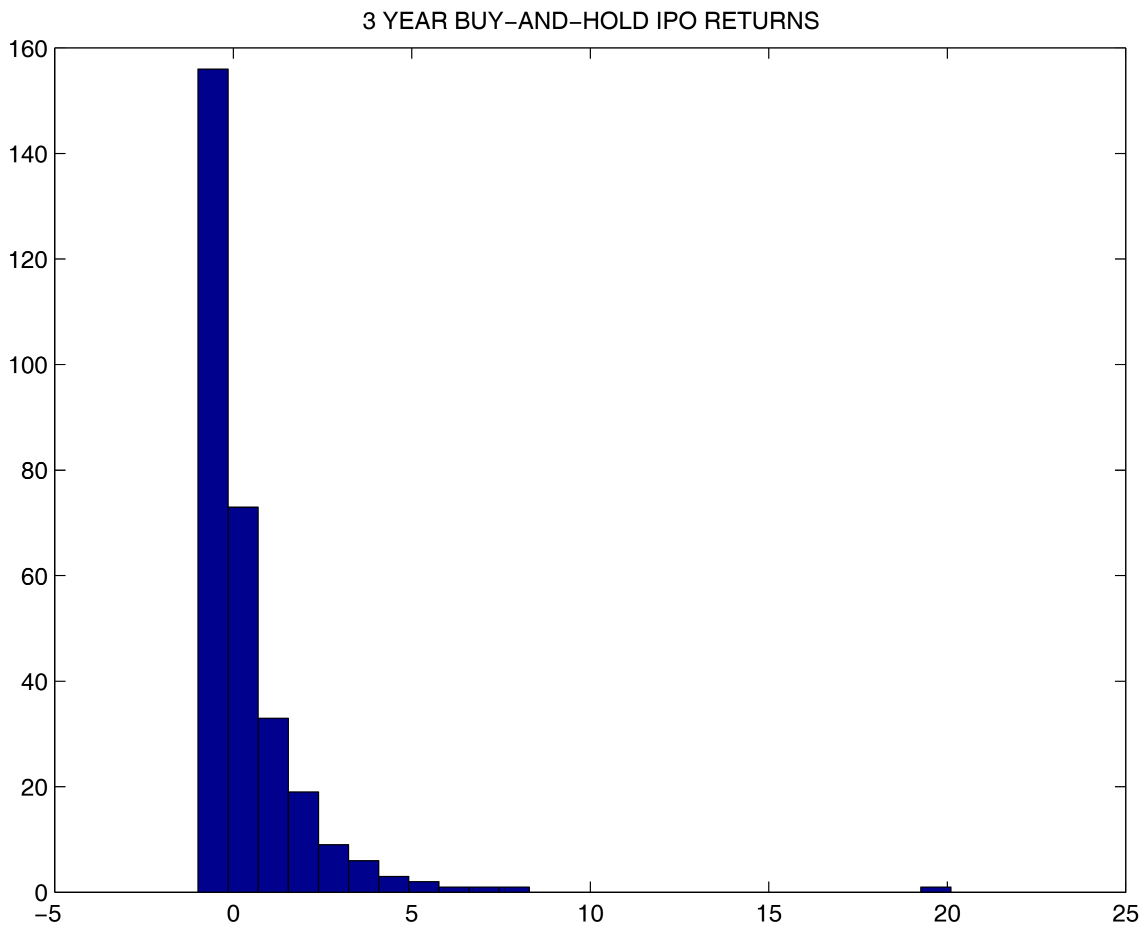

Figure 8. The histogram shows the net 3-year buy-and-hold returns for 1525 IPOs between 1975 and 1984.

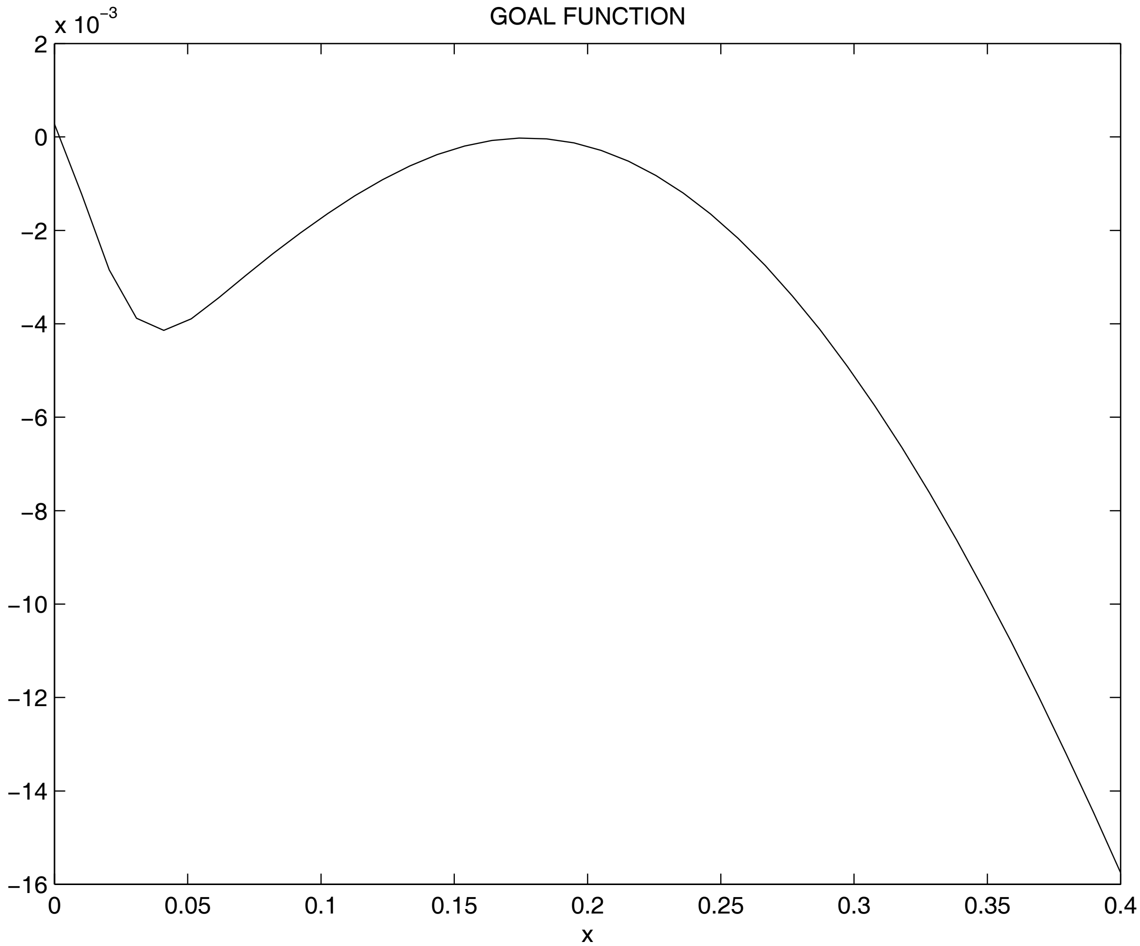

Figure 9. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x in a positively-skewed security to his current holdings of a normally distributed market portfolio. The distribution of the positively-skewed security is the historical distribution of 3-year buy-and-hold IPO returns between 1975 and 1984.